UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/808,138	03/24/2004	Samuel Fineberg	200314538-1	5258

22879 7590 07/11/2008

HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

EXAMINER
TSUI, DANIEL

ART UNIT	PAPER NUMBER
2185	

NOTIFICATION DATE	DELIVERY MODE
07/11/2008	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

JERRY.SHORMA@HP.COM
mkraft@hp.com
ipa.mail@hp.com

UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Application Number: 10/808,138
Filing Date: March 24, 2004
Appellant(s): FINEBERG ET AL.

Mark E. Scott
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed April 17, 2008 appealing from the Office action mailed February 26, 2008.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

No amendment after final has been filed.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

7,251,713	ZHANG	7-2007
6,477,617	GOLDING	11-2002
5,479,628	OLSON	12-1995
7,266,645	GARG	9-2007

ROMANOW, A., and BAILEY, S. "An Overview of RDMA Over IP" The Internet Society, 2002. Retrieved from the Internet: <URL: http://datatag.web.cern.ch/datatag/pfldnet2003/papers/romanow.pdf>.

"The Authoritative Dictionary of IEEE Standards Terms" 2000, IEEE. 7th Edition. page 46

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claims 1, 2, 8, 10, and 11 are rejected under 35 U.S.C. 103(a) as being unpatentable over Zhang (US 7,251,713) in view of RDMA ("An Overview of RDMA over IP").

As per claim 1, Zhang teaches a persistent memory access system comprising:

a primary network persistent memory unit comprising:

a network interface communicatively coupled to at least one client processor node over a communication system (storage controllers 108 provide network interface connections with clients and other controllers, see figure 1 and column 3, lines 42-45);

a primary region coupled to the network interface and configured to store information (storage unit 116, see figure 1), the primary region is assigned to a client process on a client processor node and is configured to store information received from the client process (it is an inherent property of network attached storage devices to store information from the client processor nodes; see column 5, lines 34-38);

a mirror nPMU comprising:

a network interface communicatively coupled to at least one client processor node and the primary nPMU over a communication system (storage controller 112 provide network interface connections with clients and other controllers, see figure 1 and column 3, lines 42-45);

a mirror region coupled to the network interface of the mirror nPMU and configured to store information (secondary storage 118, see figure 1; column 5, lines 7-9), the mirror region is assigned to the client process and stores the information received from the client process.

Zhang does not teach the memory units being coupled to the processor nodes over a remote direct memory access enabled communication system that executes single byte RDMA requests received through the network interface. The non-patent literature teaches the use of RDMA as a way for connecting memory units across a network (see page 2, last paragraph). It would have been obvious at the time the invention was made to a person of ordinary skill in the art for the network persistent memory units taught by Zhang to be connected over an RDMA enabled communication system so that the memory units can be directly accessed by the processes running on each client.

As per claim 2, Zhang teaches that the nPMUs are physically separate units and are characterized by separate fault domains (see column 4, lines 59-61).

As per claim 8, Zhang teaches the system further comprising a persistent memory unit library residing in the client load that comprises functions that permit

directly writing and reading information to the regions (interface software 307, see figure 2 and column 5, lines 62-65).

As per claims 10 and 11, Zhang teaches the system further comprising a persistent memory manager coupled to the processor node for creating the primary and mirror regions on the storage devices (the storage controllers 108 and 112 serve to control the storage areas 116 and 118 and allow data to be stored on the devices). This functionality would include allocating and deallocating regions for use.

Claims 3 and 4 are rejected under 35 U.S.C. 103(a) as being unpatentable over Zhang in view of RDMA and further in view of Golding (US 6,477,617).

As applied in the rejection above, the combination of Zhang and RDMA teaches a persistent memory access system with a primary region and a mirror region. The combination does not teach the memory regions comprising virtual addresses corresponding to the physicals locations where the information is stored. Golding teaches a memory storage system that uses virtual addresses so that data can be stored across multiple physical devices while still appearing to be on one storage (see column 8, lines 45-56). Golding also teaches translating between the virtual addresses to physical addresses. It would have been obvious at the time the invention was made to a person of ordinary skill in the art to use virtual addresses for both the primary storage region and the mirror region so that the data stored to these regions can be stored across multiple physical devices while appearing to be on a single unit. It would

have also been obvious to perform the address translation so that the clients using virtual addresses can access the physical locations where data is to be stored.

Claim 5 is rejected under 35 U.S.C. 103(a) as being unpatentable over Zhang in view of RDMA and Golding and further in view of Olson (US 5,479,628).

As applied in the rejection above, the combination of Zhang, RDMA, and Golding teaches a persistent memory access system that uses virtual addresses. The references do not teach using a base pointer corresponding to a difference in the primary virtual address and the corresponding client address for translating. Olson teaches a system that translates virtual to physical addresses and uses a base pointer. It would have been obvious for the claimed system to also use a pointer to perform virtual to physical address translation since it is necessary in a known technique of performing such address translations.

Claims 6 and 7 are rejected under 35 U.S.C. 103(a) as being unpatentable over Zhang in view of RDMA and further in view of Garg (US 7,266,645).

As applied in the rejection above, the combination of Zhang and RDMA teaches a persistent memory access system with a primary region and a mirror region. The combination does not teach the system comprising metadata identifying the regions assigned to the client process or caching the metadata. Garg teaches a system that uses metadata for data objects, the metadata including locations that the data is stored (see column 4, lines 8-11). Garg also teaches caching the metadata (cached metadata

90b, see column 3, lines 47-49). It would have been obvious at the time the invention was made to a person of ordinary skill in the art to use metadata to identify the regions where the data is stored and to cache the metadata so the system can access it without having to go to the storage.

Claim 9 is rejected under 35 U.S.C. 103(a) as being unpatentable over Zhang in view of RDMA and further in view of API (IEEE dictionary).

As applied in the rejection above, the combination of Zhang and RMDA teaches a persistent memory access system with a primary region and a mirror region. The combination does not teach an API residing in the client node that causes the client process to access the functions of the PMU library. However, it was well known in the art at the time the invention was made for computer systems to use APIs as the interface between applications (i.e. the client process) and the system (see definition in IEEE dictionary). Therefore it would have been obvious for the client processor node to include an API that would allow the client processes to access the library functions that perform reading and writing to the storage devices.

(10) Response to Argument

Regarding claims 1-2 and 8, Appellant argues that the combination of Zhang and RDMA is improper since Zhang is directed to a large block-level data storage devices and file-level mirroring and byte-level RDMA access would run contrary to the intended purposes. In response, Examiner points out that Zhang teaches, generally, a system

with primary and secondary storages connected over a network and mirroring across the network between the two storages. The use of large block-level data storage devices and file-level mirroring with snapshots is merely one embodiment of the Zhang system and not necessarily essential to fulfill the intended purpose of mirroring data to a remote location. Zhang is relied upon to teach the configuration and arrangement of the processors, storages, and network connections. However, as detailed in the rejections, it would have been obvious to replace the file-level mirroring and storages with storages that facilitate byte-level RDMA access that were known at the time the invention was made since doing so would have been a simple substitution with a predictable result. Although such a change would require modifications to the specific details of operation within the Zhang system, such a change would not affect the intended purpose taught by Zhang, i.e. mirroring to a remote backup site, or the principle of operation, i.e. copying data from a primary storage to a secondary storage. Such a modification would have enabled the additional benefits of byte-level RDMA such as enabling copying without direct control of the processor.

Regarding claim 10, Appellant argues that Zhang does not teach the claimed persistent memory manager since the storage controllers 108 and 112 of Zhang are relied upon to teach the claimed persistent memory units. In response, Examiner points out that Appellant has misconstrued the rejection from the Office action. Examiner has clarified the rejection to show that the primary storage 116 and secondary storage 118 are what is relied to teach the memory units. Storage controllers 108 and 112 are relied upon to provide the teaching for the network interface and the persistent memory

manager. This meets the limitations of the claims since the persistent memory manager is claimed to be coupled with the client processor.

Appellant also argues that Zhang and RDMA do not teach the persistent memory manager capable of allocating space on a non-locally coupled storage device. In response, Examiner points out that this limitation is not required in the claim since the broadest reasonable interpretation of the language could include multiple persistent memory managers coupled to local client processors so long as the primary nPMU creates a primary region and a mirror nPMU creates a mirror region. The claim does not require that one persistent memory manager perform both allocations or that both allocations be performed together with some special relationship, only that both allocations be capable of being performed. Thus Zhang teaches the limitations of the claim since it teaches a PMM capable of causing a primary nPMU to create a primary region a PMM capable of causes the mirror nPMU to create a mirror region.

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,

/Daniel Tsui/

Examiner, Art Unit 2185

Conferees:

/Sanjiv Shah/

Supervisory Patent Examiner, Art Unit 2185

/Manorama Padmanabhan/

QAS, TC2100, WG2180